

IN THE CLAIMS:

- 1 1. (Previously Presented) A method for a storage operating system implemented in a
2 storage system to optimize an amount of readahead data retrieved from
3 a data container of the storage system, comprising:
 - 4 managing a separate set of readahead metadata for each individual read stream of
5 a plurality of read streams, wherein the read stream is a set of one or more client read
6 requests to retrieve data from a contiguous range of file offsets within a requested file;
 - 7 storing the separate set of readahead metadata in a corresponding individual
8 readset data structure for the each individual read stream;
 - 9 maintaining, in the individual readset data structure, a plurality of factors for the
10 each read stream, the plurality of factors determining the amount of readahead data
11 retrieved from the data container for the one or more client read requests;
 - 12 receiving a client read request associated with a particular read stream at the
13 storage system and determining an amount of readahead data to retrieve from the data
14 container for the particular read stream;
 - 15 receiving a next client read request associated with the particular read stream;
 - 16 locating a readset data structure for the particular read stream;
 - 17 determining whether the storage operating system is permitted to retrieve
18 readahead data from the data container in response to the received next client read
19 request; and
 - 20 if it is determined that the storage operating system is permitted to retrieve
21 readahead data from the data container, performing the steps of,
 - 22 (i) modifying one or more of the plurality of factors stored within the readset data
23 structure for the particular read stream;
 - 24 (ii) adjusting, in response to modifying the one or more factors, an amount of
25 readahead data to retrieve for the next client read request associated with the particular
26 read stream, wherein the amount of readahead data to retrieve for the next client read
27 request is different from the amount of readahead data retrieved for the client read
28 request;

29 (iii) retrieving the adjusted amount of readahead data from the data container; and
30 (iv) determining if the readset data structure meets a criteria for being updated,
31 and if the readset data structure meets the criteria then updating the readset data structure.

1 2. (Original) The method of claim 1, wherein the data container is a file, directory, vdisk
2 or lun.

1 3. (Original) The method of claim 1, wherein the storage operating system is determined
2 to be permitted to retrieve readahead data from the data container when the client-
3 requested data extends the read stream past a predetermined next readahead value.

1 4. (Original) The method of claim 3, wherein the predetermined next readahead value is
2 stored in a readset data structure associated with the read stream.

1 5. (Previously Presented) The method of claim 3, wherein the predetermined next
2 readahead value is updated based on a percentage of the adjusted amount of readahead
3 data.

1 6. (Previously Presented) The method of claim 1, wherein a read-access style associated
2 with the data container is one of the plurality of factors used to select the amount of
3 readahead data.

1 7. (Previously Presented) The method of claim 6, wherein the adjusted amount of
2 readahead data equals zero if the read-access style corresponds to a random read-access
3 style.

1 8. (Previously Presented) The method of claim 1, wherein a number of client read
2 requests processed in the read stream is one of the plurality of factors used to select the
3 amount of readahead data.

- 1 9. (Original) The method of claim 8, wherein the number of client read requests
- 2 processed in the read stream is stored as a count value in a readset data structure
- 3 associated with the read stream.

- 1 10. (Previously Presented) The method of claim 1, wherein the amount of client-requested data is one of the plurality of factors used to select the amount of readahead data.

- 1 11. (Previously Presented) The method of claim 10, wherein the adjusted amount of readahead data is set equal to a predetermined upper limit for large amounts of client-requested data.

- 1 12. (Previously Presented) The method of claim 1, wherein the adjusted amount of readahead data is doubled if the number of client read requests processed in the read stream is greater than a first threshold value.

- 1 13. (Original) The method of claim 1, wherein the client-requested data is identified as read-once data when either (i) the number of client read requests processed in the read stream is greater than a second threshold value or (ii) a set of metadata associated with the read stream indicates that the client-requested data is read-once data.

- 1 14. (Previously Presented) The method of claim 1, wherein the adjusted amount of readahead data is stored in one or more buffers enqueued on a flush queue, the flush queue being configured to reuse buffers after a predetermined period of time.

- 1 15. (Original) The method of claim 14, wherein the predetermined period of time equals two seconds.

1 16. (Previously Presented) An apparatus configured to implement a storage operating
2 system that optimizes an amount of readahead data retrieved from a data container of the
3 apparatus, comprising:

4 means for managing a separate set of readahead metadata for each individual read
5 stream of a plurality of read streams, wherein the read stream is a set of one or more
6 client read requests to retrieve data from a contiguous range of file offsets within a
7 requested file;

8 means for storing the separate set of readahead metadata in a corresponding
9 individual readset data structure for the each individual read stream;

10 means for maintaining, in the individual readset data structure, a plurality of
11 factors for the each read stream, the plurality of factors determining the amount of
12 readahead data retrieved from the data container for the one or more client read requests;

13 means for receiving a client read request associated with a particular read stream
14 at the storage system and determining an amount of readahead data to retrieve from the
15 data container for the particular read stream;

16 means for receiving a next client read request associated with the particular read
17 stream;

18 means for locating a readset data structure for the particular read stream;

19 means for determining whether the storage operating system is permitted to
20 retrieve readahead data from the data container in response to the received next client
21 read request; and

22 if it is determined that the storage operating system is permitted to retrieve
23 readahead data from the data container,

24 (i) means for modifying one or more of the plurality of factors stored within the
25 readset data structure for the particular read stream;

26 (ii) means for adjusting, in response to modifying the one or more factors, an
27 amount of readahead data to retrieve for the next client read request associated with the
28 particular read stream, wherein the amount of readahead data to retrieve for the next
29 client read request is different from the amount of readahead data retrieved for the client
30 read request;

31 (iii) means for retrieving the adjusted amount of readahead data from the data
32 container; and
33 (iv) means for determining if the readset data structure meets a criteria for being
34 updated, and if the readset data structure meets the criteria then updating the readset data
35 structure.

1 17. (Original) The apparatus of claim 16, wherein the data container is a file, directory,
2 vdisk or lun.

1 18. (Original) The apparatus of claim 16, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 19. (Previously Presented) The apparatus of claim 18, further comprising means for
2 updating the predetermined next readahead value based on a percentage of the adjusted
3 amount of readahead data.

1 20. (Previously Presented) The apparatus of claim 16, wherein the plurality of factors
2 used to select the amount of readahead data ~~includes~~ comprises at least one of:
3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 21. (Previously Presented) The apparatus of claim 16, wherein the adjusted amount of
2 readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

1 22. (Previously Presented) A storage system configured to optimize an amount of
2 readahead data retrieved from a data container of the storage system, ~~the storage system~~
3 comprising:
4 a network adapter for receiving a client read request for a particular read stream;
5 a memory configured to store instructions for implementing a storage operating
6 system that performs the steps of:
7 managing a separate set of readahead metadata for each individual read stream of
8 a plurality of read streams, wherein the read stream is a set of one or more client read
9 requests to retrieve data from a contiguous range of file offsets within a requested file;
10 storing the separate set of readahead metadata in a corresponding individual
11 readset data structure for the each individual read stream;
12 maintaining, in the individual readset data structure, a plurality of factors for the
13 each read stream, the plurality of factors determining the amount of readahead data
14 retrieved from the data container for the one or more client read requests;
15 receiving a client read request associated with a particular read stream at the
16 storage system and determining an amount of readahead data to retrieve from the data
17 container for the particular read stream;
18 receiving a next client read request associated with the particular read stream;
19 locating a readset data structure for the particular read stream;
20 determining whether the storage operating system is permitted to retrieve
21 readahead data from the data container in response to the received next client read
22 request; and
23 if it is determined that the storage operating system is permitted to retrieve
24 readahead data from the data container, performing the steps of,
25 (i) modifying one or more of the plurality of factors stored within the readset data
26 structure for the particular read stream;
27 (ii) adjusting, in response to modifying the one or more factors, an amount of
28 readahead data to retrieve for the next client read request associated with the particular
29 read stream, wherein the amount of readahead data to retrieve for the next client read

30 request is different from the amount of readahead data retrieved for the client read
31 request;
32 (iii) retrieving the adjusted amount of readahead data from the data container; and
33 (iv) determining if the readset data structure meets a criteria for being updated,
34 and if the readset data structure meets the criteria then updating the readset data structure.

1 23. (Original) The storage system of claim 22, wherein the data container is a file,
2 directory, vdisk or lun.

1 24. (Original) The storage system of claim 22, wherein the storage operating system is
2 determined to be permitted to retrieve readahead data from the data container when the
3 client-requested data extends the read stream past a predetermined next readahead value.

1 25. (Previously Presented) The storage system of claim 24, wherein the predetermined
2 next readahead value is updated based on a percentage of the adjusted amount of
3 readahead data.

1 26. (Previously Presented) The storage system of claim 22, wherein the plurality of
2 factors used to select the amount of readahead data includes comprises at least one of:
3 (i) the amount of client-requested data,
4 (ii) a number of client read requests processed in the read stream, and
5 (iii) a read-access style associated with the data container.

1 27. (Previously Presented) The storage system of claim 22, wherein the adjusted amount
2 of readahead data is doubled if the number of client read requests processed in the read
3 stream is greater than a first threshold value.

1 28. (Previously Presented) A computer-readable media containing program instructions
2 executed by a processor, comprising:

3 program instructions that manage a separate set of readahead metadata for each
4 individual read stream of a plurality of read streams, wherein the read stream is a set of
5 one or more client read requests to retrieve data from a contiguous range of file offsets
6 within a requested file;

7 program instructions that store the separate set of readahead metadata in a
8 corresponding individual readset data structure for the each individual read stream;

9 program instructions that maintain, in the individual readset data structure, a
10 plurality of factors for the each read stream, the plurality of factors determining the
11 amount of readahead data retrieved from the data container for the one or more client
12 read requests;

13 program instructions that receive a client read request associated with a particular
14 read stream at the storage system and determining an amount of readahead data to
15 retrieve from the data container for the particular read stream;

16 program instructions that receive a next client read request associated with the
17 particular read stream;

18 program instructions that locate a readset data structure for the particular read
19 stream;

20 program instructions that determine whether the storage operating system is
21 permitted to retrieve readahead data from the data container in response to the received
22 next client read request; and

23 if it is determined that the storage operating system is permitted to retrieve
24 readahead data from the data container, program instructions that,

25 (i) modify one or more of the plurality of factors stored within the readset data
26 structure for the particular read stream;

27 (ii) adjust, in response to modifying the one or more factors, an amount of
28 readahead data to retrieve for the next client read request associated with the particular
29 read stream, wherein the amount of readahead data to retrieve for the next client read
30 request is different from the amount of readahead data retrieved for the client read
31 request;

32 (iii) retrieve the adjusted amount of readahead data from the data container; and

33 (iv) determine if the readset data structure meets a criteria for being updated, and
34 if the readset data structure meets the criteria then update the readset data structure.

1 29. (Original) The computer-readable media of claim 28, wherein the data container is a
2 file, directory, vdisk or lun.

1 30. (Previously Presented) The method of claim 1, wherein the retrieved readahead data
2 is stored in one or more buffers, the one or more buffers ~~containing~~ comprising a flush
3 queue, the flush queue being configured to reuse buffers after a predetermined period of
4 time.

1 31. (Previously Presented) The method of claim 30, wherein the read stream
2 corresponds to a read-once data transfer and data retrieved from the data container is
3 stored in the flush queue.

1 32. (Previously Presented) The method of claim 30, wherein the retrieved readahead
2 data is stored in the flush queue.

1 33. (Previously Presented) The method of claim 30, wherein one or more buffers
2 accessed from the flush queue are re-enqueued on a normal queue.

1 34-75. (Cancelled).